Mon. Not. R. Astron. Soc. **000**, 1–9 (2016) Printed March 16, 2017 (MN LATEX style file v2.2)

Quantifying lost information due to covariance matrix estimation in parameter inference

Elena Sellentin, Alan F. Heavens
Imperial Centre for Inference and Cosmology (ICIC), Department of Physics, Imperial College, Blackett Laboratory, Prince Consort Road, London SW7 2AZ, U.K.

Accepted 17th October 2016. Received 17th October 2016; in original form 6th September 2016

arXiv:1609.00504v2 [astro-ph.CO] 15 Mar 2017

ABSTRACT

Parameter inference with an estimated covariance matrix systematically loses information due to the remaining uncertainty of the covariance matrix. Here, we quantify this loss of precision and develop a framework to hypothetically restore it, which allows to judge how far away a given analysis is from the ideal case of a known covariance matrix. We point out that it is insufficient to estimate this loss by debiasing a Fisher matrix as previously done, due to a fundamental inequality that describes how biases arise in non-linear functions. We therefore develop direct estimators for parameter credibility contours and the figure of merit, finding that significantly fewer simulations than previously thought are sufficient to reach satisfactory precisions. We apply our results to DES Science Verification weak lensing data, detecting a 10% loss of information that increases their credibility contours. No significant loss of information is found for KiDS. For a Euclid-like survey, with about 10 nuisance parameters we find that 2900 simulations are sufficient to limit the systematically lost information to 1%, with an additional uncertainty of about 2%. Without any nuisance parameters 1900 simulations are sufficient to only lose 1% of information. We further derive estimators for all quantities needed for forecasting with estimated covariance matrices. Our formalism allows to determine the sweetspot between running sophisticated simulations to reduce the number of nuisance parameters, and running as many fast simulations as possible.

Key words: methods: data analysis – methods: statistical – cosmology: observations

1 INTRODUCTION

For surveys of the cosmic large-scale structure as for example in Abbott et al. (2015); Heymans & et al. (2013); Laureijs et al. (2011); Jain et al. (2015), cosmology currently finds itself in the situation of having difficulties reliably describing measurement uncertainties via a covariance matrix. Whereas analytical approximations for the covariance matrix exist and have been exploited e.g. in KiDS (Hildebrandt et al. 2016), covariance matrices are usually estimated from numerical simulations with the hope that these model the specifics of a survey and the non-linear structure growth better than analytical approximations (Harnois-Déraps & van Waerbeke 2015; Blot et al. 2015; Sato et al. 2011).

Uncertainty in the covariance matrix inevitably leads to loss of information, and a modified likelihood function (Sellentin & Heavens 2016). Here we investigate further the loss of information, to quantify the expected gains or losses obtainable by increasing or decreasing the number of simulated datasets.

We begin with an unbiased sample estimator for the

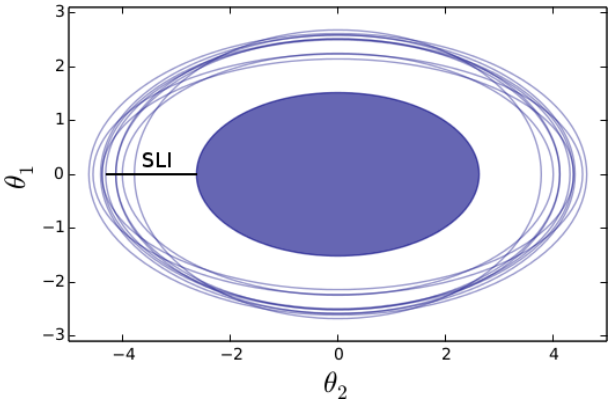

Figure 1. Toy model, estimating linear parameters from a t-distribution with multiple Wishart-sampled covariance matrices (open contours), or from a Gaussian distribution with the true covariance matrix (solid contour). The gap between the solid contour and the open contours illustrates the systematically lost information (SLI) due to the uncertain covariance matrix S. This gap diminishes when increasing the number of simulations N_s.

covariance matrix. If N_s independent simulations each yield a synthetic data vector $\boldsymbol{X}_i$, with the sample average being $\bar{\boldsymbol{X}} = \frac{1}{N_s}\sum_{i=1}^{N_s}\boldsymbol{X}_i$, then an unbiased estimator of the true but unknown covariance matrix $\boldsymbol{\Sigma}$ is

$$\mathsf{S} = \frac{1}{N_s-1}\sum_{i=1}^{N_s}(\boldsymbol{X}_i-\bar{\boldsymbol{X}})(\boldsymbol{X}_i-\bar{\boldsymbol{X}})^T. \quad (1)$$

This estimator is a matrix-variate random variable and in Sellentin & Heavens (2016), we propagated the uncertainty on the true covariance $\boldsymbol{\Sigma}$ into parameter inference, by transferring the randomness of S onto $\boldsymbol{\Sigma}$ by virtue of a prior and Bayes' theorem. This lead to the likelihood of a p-dimensional dataset $\boldsymbol{X}_{\rm o}$, conditioned on the mean $\boldsymbol{\mu}$ and S from N_s simulations:

$$P(\boldsymbol{X}_{\rm o}|\boldsymbol{\mu},\mathsf{S},N_s) = \frac{\bar{c}_p|\mathsf{S}|^{-1/2}}{\left[1+\frac{Q}{N_s-1}\right]^{\frac{N_s}{2}}}. \quad (2)$$

This is an ellipsoidally contoured t-distribution, using

$$Q = (\boldsymbol{X}_{\rm o}-\boldsymbol{\mu})^T\mathsf{S}^{-1}(\boldsymbol{X}_{\rm o}-\boldsymbol{\mu}), \quad (3)$$

and the normalization

$$\bar{c}_p = \frac{\Gamma\left(\frac{N_s}{2}\right)}{[\pi(N_s-1)]^{p/2}\,\Gamma\left(\frac{N_s-p}{2}\right)}. \quad (4)$$

Previous approximate results had been proposed in a series of papers (Hartlap et al. 2007; Dodelson & Schneider 2013; Taylor et al. 2013; Taylor & Joachimi 2014; Percival et al. 2014) which employ only the first two moments instead of the entire distribution of S, by construction maintaining a Gaussian likelihood. In comparison to a Gaussian, a t-distribution has broader wings, and a more peaked core. Its width is primarily set by the number of simulations, as these determine how certain our estimate of the true covariance matrix is. The lower the number of simulations, the wider is the t-distribution which reflects the systematic loss of information due to estimating the covariance matrix. The obvious question in this framework is therefore how much information is lost for a certain finite N_s, in comparison to a known covariance matrix and this is the question we address in this paper.

We illustrate this loss with a toy model in Fig. 1 for a survey with similar characteristics to ESA's Euclid weak lensing survey (Laureijs et al. 2011). Our agnostic model for Euclid assumes 10 redshift bins, leading to approximately $p = 1500$ data points. For Fig. 1, we use $N_s = 1550$ simulations and $N_p = 50$ parameters that include cosmological parameters and a representative number of nuisance parameters. The relative uncertainties due to estimated covariance matrices do not depend on any further physical specifications. The solid contour is the joint 1σ credible region of two arbitrary parameters θ_1 and θ_2, derived from an assumed true covariance matrix. In contrast, the open contours are derived from different draws of estimated covariance matrices, and are systematically inflated, due to losing information when estimating the covariance matrix.

We will assess this loss by developing a Fisher matrix approximation to the t-distribution in order to completely include the uncertainty of $\boldsymbol{\Sigma}$. However, we will not stop at the Fisher matrix level, as previous works have done, as an important inequality enforces that the lost information of joint credibility regions and Figures of Merit (FoM) cannot be calculated from the information lost in a Fisher matrix: Jensen's inequality describes the fundamental problem that non-linear functions and averages of estimators do not commute. For any matrix variate function f of the estimated covariance matrix S, Jensen's inequality reads

$$f(\langle\mathsf{S}\rangle) \leqslant \langle f(\mathsf{S})\rangle \quad (5)$$

for convex functions f, and the inequality inverts for concave functions. Here, angular brackets denote averaging, and the equality holds for linear functions only. This inequality implies that the mean loss of information for credible regions of jointly estimated parameters cannot be assessed by debiasing a Fisher matrix, because joint credibility regions are non-linear functions of the Fisher matrix. Likewise, the mean loss of information in the figure of merit cannot be predicted from the loss of information in the Fisher matrix, again due to the non-linear dependence between the two. We find it therefore necessary to extend the works of Dodelson & Schneider (2013); Taylor et al. (2013); Taylor & Joachimi (2014), which focused on Fisher matrices; here, we calculate the loss of information directly from the FoM and parameter credibility contours. We will find that the number of simulations needed to acquire a certain precision in the FoM or parameter credibility contours (or related quantities) depends on the estimator, but is generally lower than predicted by previous forecasts. This is mainly because we will use the more accurate t-distribution Eq. (2) which was not available to previous forecasts.

We establish the mathematical framework in Sects. 2-4 and then validate it by a full Monte Carlo Markov Chain run in Sect. 5. Conclusions for the information lost in current and future surveys are given in Sect. 5 and the discussion at the end of the paper.

2 FISHER MATRIX

We begin by deriving the Fisher matrix of the t-distribution Eq. (2). We assume that only the mean $\boldsymbol{\mu}$ depends on N_p parameters $\theta_1 \ldots \theta_{N_p}$, but the covariance matrix is parameter independent. This is strictly not the case, but is almost universally assumed, and is an excellent approximation as in typical cases the extra information from the parameter dependence of the covariance matrix is very small (Tegmark et al. 1997).

We employ Einstein's summation convention and let roman indices run over the data, $i \in [1,p]$ and greek indices over the parameters, $\alpha \in [1,N_p]$. The Fisher matrix is then

$$\mathsf{F}(\mathsf{S})_{\alpha\beta} = \langle(\ln P)_{,\alpha}(\ln P)_{,\beta}\rangle \quad (6)$$

where commas denote partial derivatives and the angular brackets denote the marginalization over the data by integrating over the t-distribution, with S kept fixed. Note that $\mathsf{F}(\mathsf{S})$ is a random object, since S is. Evaluating the gradients of Eq. (2), we have

$$\mathsf{F}_{\alpha\beta}(\mathsf{S}) = \left(\frac{N_s}{N_s-1}\right)^2 \bar{c}_p|S|^{-\frac{1}{2}}S^{-1}_{ij}S^{-1}_{kl}\mu_{j,\alpha}\mu_{l,\beta}\,I_{ik}, \quad (7)$$

where we defined the integral

$$I_{ik} \equiv \int \mathrm{d}^pX\,\frac{1}{\left(1+\frac{Q}{N_s-1}\right)^{N_s/2+2}}\,(\boldsymbol{X}-\boldsymbol{\mu})_i(\boldsymbol{X}-\boldsymbol{\mu})_k. \quad (8)$$

To solve the integral, we isotropize the data by diagonalising and whitening S. Using an orthogonal rotation matrix R, $\mathsf{S}^{-1} = \mathsf{R}^\mathsf{T}\Lambda\mathsf{R}$, where Λ is diagonal. Hence

$$Q = z_j \Lambda_{jk} z_k \tag{9}$$

where $z_i = R_{ij}(\boldsymbol{X} - \boldsymbol{\mu})_j$ and the Jacobian from X_i to z_i is unity. Suspending the summation convention, we whiten the covariance by rescaling $Z_i = \Lambda_{ii}^{\frac{1}{2}} z_i$, leading to

$$I_{ik} = |S|^{\frac{1}{2}} \sum_{st} \frac{R_{si} R_{tk}}{\Lambda_{tt}^{\frac{1}{2}} \Lambda_{ss}^{\frac{1}{2}}} J_{st}, \tag{10}$$

where the integral is

$$J_{st} = \int \mathrm{d}^p Z \frac{Z_s Z_t}{\left[1 + \frac{1}{N_s - 1} \sum_q Z_q^2\right]^{N_s/2+2}}. \tag{11}$$

By symmetry, we now have $J_{st} = 0$ unless $s = t$, in which case the integral depends only on $r^2 \equiv \sum_q Z_q^2$. We therefore transform to polar coordinates,

$$\mathrm{d}^p Z = r^{p-1} \mathrm{d}r \sin^{p-2}\phi_1 \, \sin^{p-3}\phi_2 \ldots \sin\phi_{p-1} \mathrm{d}\phi_1 \ldots \mathrm{d}\phi_p. \tag{12}$$

The angular part of the integral is A/p, where

$$A = \frac{2\pi^{p/2}}{\Gamma\left(\frac{p}{2}\right)}, \tag{13}$$

is the area of the unit sphere in the p-dimensional space of data points. The integral over r is

$$\int_0^\infty \frac{r^{p+1}\mathrm{d}r}{\left(1 + \frac{r^2}{N_s - 1}\right)^{\frac{N_s}{2}+2}} = \frac{\Gamma\left(\frac{N_s - p + 2}{2}\right)\Gamma\left(\frac{p}{2} + 1\right)}{2(N_s - 1)^{-(\frac{p}{2}+1)}\Gamma\left(\frac{N_s}{2} + 2\right)}. \tag{14}$$

The gamma functions simplify when combined with the normalization $\bar{c}_p$ from Eq. (4), giving

$$\mathsf{F}_{\alpha\beta}(\mathsf{S}) = \frac{N_s(N_s - p)}{(N_s - 1)(N_s + 2)} \sum_{ijkls} \frac{R_{si} R_{sk}}{\Lambda_{ss}} S_{ij}^{-1} S_{kl}^{-1} \mu_{j,\alpha} \mu_{l,\beta}, \tag{15}$$

where the sum over the R and Λ terms is S_{ik}. Introducing the $(p \times N_p)$ matrix A

$$\mathsf{A} = \begin{pmatrix} \boldsymbol{\mu}_{,1}^T \\ \vdots \\ \boldsymbol{\mu}_{,N_p}^T \end{pmatrix} \tag{16}$$

whose row vectors are the N_p derivatives of the mean, the Fisher matrix of the t-distribution is therefore

$$\mathsf{F}(\mathsf{S}) = \frac{N_s(N_s - p)}{(N_s - 1)(N_s + 2)} \mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T. \tag{17}$$

This Fisher matrix describes the average curvature of t-distribution around its peak, if different data sets $\boldsymbol{X}$ are drawn. It keeps the covariance matrix S fixed and is an approximation to the width of the t-distribution. Therefore, it quantifies the lost information due to S being estimated from a finite number of simulations.

A completely different use of this Fisher matrix would be to construct from it an estimator of the unknown *true* Fisher matrix, $\mathsf{F}(\Sigma) = \mathsf{A}\Sigma^{-1}\mathsf{A}^T$ with the aim of undoing the lost information in the estimation process. In a real data analysis, such lost information can only be captured by increasing the number of simulations. One should therefore not attribute any other meaning to such a Fisher matrix than being the hypothetically optimal case that could be achieved for $N_s \to \infty$.

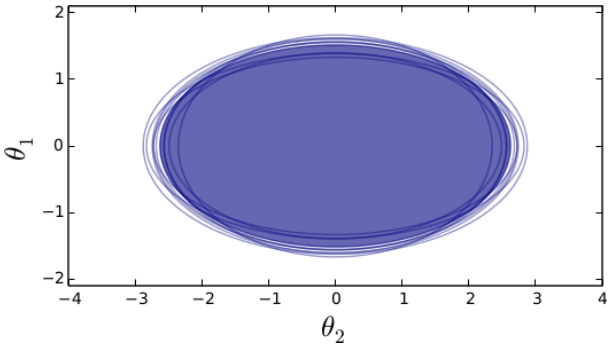

Figure 2. Scaling the confidence contours with Eq. (40) reduces their size to the optimal case of a known covariance matrix. The remaining scatter is then described by Eq. (38) and depicted in Fig. 3 for a Euclid like survey.

In a Bayesian way, estimators of unknown quantities would ideally be constructed by marginalizing over their full distribution, conditional on all available information. In this paper, we will however soon deal with non-linear functions of estimators, where their full distribution becomes analytically intractable. Fortunately, a distribution can be written in terms of its moments, provided that they are finite. We therefore employ a moment-wise description of various estimators, as it is analytically tractable. By sampling the random processes under consideration, we furthermore find that working with the first two moments of these distributions already captures the essential information that we seek.

We now estimate the Fisher matrix for Σ, given that we only know S. Since S^{-1} follows an inverse Wishart distribution, the distribution of $\mathsf{F}(\mathsf{S})$ is also inverse Wishart, with (Gupta & Nagar 2000)

$$\mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T \sim \mathcal{W}_{N_p}^{-1}\left(n - p + N_p, (\mathsf{A}\Sigma^{-1} n \mathsf{A}^T)\right). \tag{18}$$

where $n = N_s - 1$. The mean Fisher matrix for a t-distribution, when also averaged over different samples of S, is therefore

$$\langle \mathsf{F}(\mathsf{S}) \rangle_\mathsf{S} = \frac{N_s(N_s - p)}{(N_s - 1)(N_s + 2)} \langle \mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T \rangle_\mathsf{S}. \tag{19}$$

where by property of the inverse Wishart distribution

$$\langle \mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T \rangle = \frac{n}{n - p - 1} (\mathsf{A}\Sigma^{-1}\mathsf{A}^T), \tag{20}$$

where the number of data points p appears because it is the dimension of the data covariance matrix Σ. Our interest lies in estimating the true $\mathsf{A}\Sigma^{-1}\mathsf{A}^T$. Consequently, we have

$$\langle \mathsf{F}(\mathsf{S}) \rangle_\mathsf{S} = \frac{N_s(N_s - p)}{(N_s + 2)(N_s - p - 2)} (\mathsf{A}\Sigma^{-1}\mathsf{A}^T). \tag{21}$$

We therefore see that the Fisher matrix of the estimator S (random through its dependence on S) is biased with respect to the true Fisher matrix $\mathsf{A}\Sigma^{-1}\mathsf{A}^T$. This bias is however not an unexpected nuisance, rather it describes the systematic loss of information on the parameters, due to estimating S from a finite number of simulations. We therefore often refer to the bias as systematically lost information (SLI) when

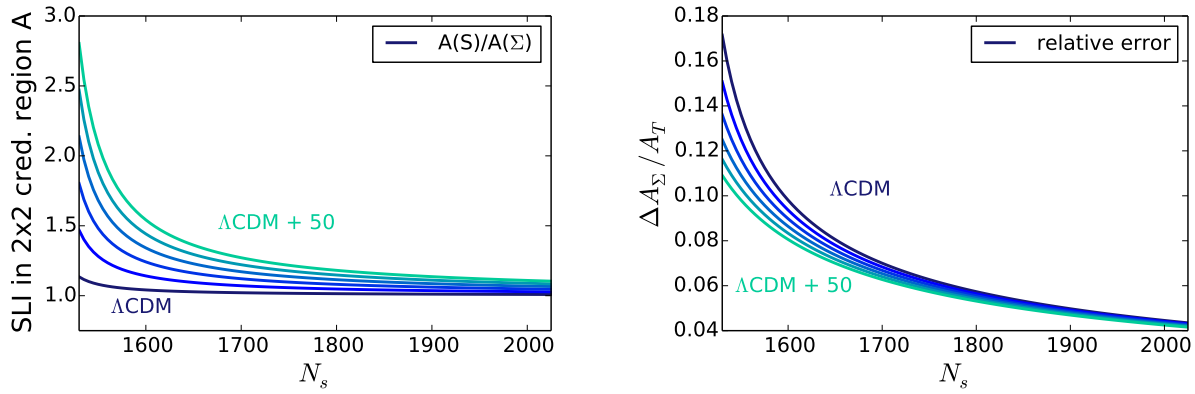

Figure 3. Left: systematically lost information (SLI) in each 2×2 parameter credibility contour, as e.g. displayed in a triangle plot, when a total of $N_p > 2$ parameters are estimated. The SLI describes the factor by which the credibility contours are systematically inflated, see also Fig. 1. Right: the statistical uncertainty of forecasting the optimally achievable credibility contours, when extrapolating from an initial estimator of N_s simulations via Eq. (38). This produces the open contours of Fig. 2 and the quantity here plotted corresponds to their observed scatter in Fig. 2 around the true contour. This plot assumes a Euclid-like number of data points, $p = 1500$. Starting at $N_p = 6$ for the ΛCDM line, the other lines subsequently add nuisance parameters in steps of 10.

emphasising its physical meaning. Removing this bias corresponds to a hypothetical restoration of the lost information, and a reduction of the bias can be achieved by increasing N_s. Eq. (21) allows a cost-benefit analysis to determine whether this is worth doing.

We therefore divide the Fisher matrix Eq. (17) by the prefactor of Eq. (21) in order to mimic restoring the lost information, leading to

$$\mathsf{F}_\Sigma(\mathsf{S}) = \frac{N_s - p - 2}{N_s - 1}\mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T. \tag{22}$$

We keep the subscript Σ to denote that this Fisher matrix is now an estimator that extrapolates from the known S to the Fisher matrix for the optimal case where $\mathbf{\Sigma}$ is known. The main difference between Eq. (17) and Eq. (22) is that $\mathsf{F}(\mathsf{S})$ includes the SLI due to uncertainty of the estimated covariance matrix, i.e. it approximates the broad t-distribution, whereas $\mathsf{F}_\Sigma(\mathsf{S})$ is constructed to *restore* the SLI, because it is an estimator of the true Fisher matrix. $\mathsf{F}_\Sigma(\mathsf{S})$ will approximate the narrower Gaussian distribution of the parameters. Loosely speaking, $\mathsf{F}(\mathsf{S})$ approximates the outer open contours, and $\mathsf{F}_\Sigma(\mathsf{S})$ approximates the inner solid contour in Fig. 1, where the caveat is that Eq. (22) corresponds only to a 'debiased' estimator of the true Fisher matrix itself, or any linear function of it. It therefore restores lost information only in the Fisher matrix itself. Unbiased parameter confidence contours, figures of merit and inverse Fisher matrices cannot be estimated from it, as these are non-linear functions of the Fisher matrix. Biases in these will reappear due to Jensen's inequality, even if $\mathsf{F}_\Sigma(\mathsf{S})$ itself is debiased. As the bias parametrizes the lost information in this context, we therefore proceed to calculate the SLI of these quantities explicitly.

2.1 Inverse Fisher matrices

The inverse of the Fisher matrix approximates the mean parameter covariance matrix, through the inverse of Eq. (21),

$$(\langle\mathsf{F}\rangle_\mathsf{S})^{-1} = \frac{(N_s+2)(N_s-p-2)}{N_s(N_s-p)}(\mathsf{A}\mathbf{\Sigma}^{-1}\mathsf{A}^T)^{-1}, \tag{23}$$

whose bias is independent of the number of parameters N_p, as opposed to the parameter covariance matrix found by Taylor & Joachimi (2014). Further, the bias factor of Eq. (23) cannot diverge, as the Wishart distribution requires $N_s > p$ anyway. An estimator that restores the SLI for the inverse of the average peak curvature, is therefore

$$\mathsf{F}^{-1}_{\Sigma,\mathrm{peak}}(\mathsf{S}) = \frac{(N_s-1)}{(N_s-p-2)}(\mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T)^{-1}, \tag{24}$$

which is indeed the inverse of the Fisher matrix Eq. (22).

However, a second estimator for the true inverse Fisher matrix can be formed by taking the matrix inverse of the conditional Eq. (17) directly,

$$\mathsf{F}^{-1}(\mathsf{S}) = \frac{(N_s-1)(N_s+2)}{N_s(N_s-p)}(\mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T)^{-1}. \tag{25}$$

Again, this inverse Fisher matrix is random and follows a Wishart distribution (Gupta & Nagar 2000):

$$(\mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T)^{-1} \sim \mathcal{W}_{N_p}\left(n-p+N_p, (\mathsf{A}\mathbf{\Sigma}^{-1}n\mathsf{A}^T)^{-1}\right). \tag{26}$$

Its mean is therefore

$$\langle\mathsf{F}^{-1}(\mathsf{S})\rangle = \frac{(N_s-p+N_p-1)(N_s+2)}{N_s(N_s-p)}(\mathsf{A}\mathbf{\Sigma}^{-1}\mathsf{A}^T)^{-1}, \tag{27}$$

whose bias depends on the number of parameters, similar to the estimator found by Taylor & Joachimi (2014). Its equivalent that restores SLI in a single estimated covariance matrix from Eq. (25) is then

$$\mathsf{F}^{-1}_\Sigma(\mathsf{S}) = \frac{(N_s-1)}{(N_s-p+N_p-1)}(\mathsf{A}\mathsf{S}^{-1}\mathsf{A}^T)^{-1}. \tag{28}$$

The estimators Eq. (23) and Eq. (27) differ numerically because averaging and inverting are mathematical operations that do not commute. Physically, the estimators Eq. (23) and Eq. (27) correspond to different quantities, and it depends on the scientific context which of the two is of interest. Eq. (23) is the inverse of the average curvature of the t-distribution around its peak, marginalized over data $\boldsymbol{X}$ and over covariance estimators S which set the width of a t-distribution. Eq. (27) describes how the inverse Fisher

matrix for a single S scatters if the parameter inference, conditional on an estimator $\mathsf{S}_1 \sim \mathcal{W}_p(n, \boldsymbol{\Sigma})$ were repeated for other samples $\mathsf{S}_i \sim \mathcal{W}_p(n, \boldsymbol{\Sigma})$. Eq. (27) then describes that if one calculates an inverse Fisher matrix for each sample S_i individually, the mean will be biased due to the lost information. As a side note, our inverse Fisher matrix Eq. (27) differs from the estimator given in Taylor & Joachimi (2014) because we work with the Bayesian t-distribution, whereas Taylor & Joachimi (2014) use a systematically inflated Gaussian likelihood following Hartlap et al. (2007). Our main difference is however that we regard the Fisher matrix only as an intermediate step, needed to derive the quantities presented in the upcoming sections.

3 SLI IN PARAMETER CREDIBLE REGIONS

Typically, a data analysis has just a single estimated covariance matrix S at its disposal. Ultimately, our interest in it is to calculate parameter credible regions. For deriving the SLI of credible regions and the remaining scatter after hypothetically restoring it, the estimator Eq. (25) for the inverse Fisher matrix is therefore the correct choice. Since the inverse Fisher matrix is random, its eigenvectors will scatter, leading to a reorientation of the credible regions that are derived from it. Additionally, the volume of credible regions will vary. Here we concentrate on the latter.

We employ a geometrical interpretation of covariance matrices. The Fisher matrix approximation is an ellipsoidally-contoured likelihood with

$$L(\boldsymbol{\theta}) \approx \mathcal{C} \exp\left[-\frac{1}{2}(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})^T \mathsf{F}(\mathsf{S})(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})\right], \tag{29}$$

where $\mathcal{C}$ is a normalization constant and $\hat{\boldsymbol{\theta}}$ are the best-fitting parameters. The quadratic form

$$(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}})^T \mathsf{F}(\mathsf{S})(\boldsymbol{\theta}-\hat{\boldsymbol{\theta}}) = \chi^2_r(\alpha), \tag{30}$$

describes hyper-ellipsoids, centered on $\hat{\boldsymbol{\theta}}$, encasing $1-\alpha$ of the total probability. Their axes are $\boldsymbol{a}_i^{\pm} = \pm\sqrt{\chi^2_r(\alpha)\lambda_i}\boldsymbol{e}_i$, where λ_i is the ith eigenvalue of F^{-1}, and $\boldsymbol{e}_i$ is the corresponding unit-length eigenvector. The degrees of freedom r of the χ^2-distribution are the number of parameters to be estimated.

The volume of a r-dimensional unit-sphere is

$$V(r) = \frac{\pi^{r/2}}{\Gamma(\frac{r}{2}+1)}. \tag{31}$$

The inverse Fisher matrix then maps this unit hyper-sphere onto the hyper-ellipsoids that describe the parameter confidence levels. The volume of the mapped and distorted hyper-sphere then changes by the determinant of the inverse Fisher matrix, such that in r dimensions, the volume of an hyper-ellipsoid of parameter credible regions is

$$V_e(r) = \frac{\pi^{r/2}}{\Gamma(\frac{r}{2}+1)}\left[\chi^2_r(\alpha)\right]^{\frac{r}{2}}\sqrt{|\mathsf{F}^{-1}_{r\times r}|}. \tag{32}$$

If we want to measure N_p parameters then $|\mathsf{F}^{-1}_{N_p\times N_p}|$ is needed to describe the total parameter volume, whereas if we are only interested in an $r\times r$ submatrix of $\mathsf{F}^{-1}_{N_p\times N_p}$, then using its determinant in Eq. (32) will describe the parameter volume of the remaining r parameters, where the other parameters have been marginalized, due to discarding rows and columns from the inverse Fisher matrix. Consequently, for assessing the SLI of joint credible regions, we require the bias and variance of submatrices of F^{-1}, and of their determinants.

Fortunately, $\mathsf{F}^{-1}(\mathsf{S})$ is Wishart distributed, and $r\times r$ submatrices from the diagonal of a Wishart matrix are again Wishart distributed with the same degrees of freedom but dimension r (Anderson 2003). Defining the pre-factor of Eq. (17) as $\lambda = N_s(N_s-p)/[(N_s-1)(N_s+2)]$, we therefore have that an $r\times r$ submatrix of F^{-1} is distributed according to

$$\left(\lambda\mathsf{F}^{-1}(\mathsf{S})\right)_{r\times r} \sim \mathcal{W}_r(n-p+N_p, (\mathsf{A}\boldsymbol{\Sigma}^{-1}n\mathsf{A}^T)^{-1}_{r\times r}). \tag{33}$$

Further, the hth moments of the distribution of a Wishart matrix's determinant are known to be (Gupta & Nagar 2000)

$$\langle|\mathsf{F}^{-1}_{r\times r}(\mathsf{S})|^h\rangle = |(\mathsf{A}\boldsymbol{\Sigma}^{-1}\mathsf{A}^T)^{-1}_{r\times r}|^h B_r(h), \tag{34}$$

where the Fisher matrix derived from the true covariance matrix is $\mathsf{F}(\boldsymbol{\Sigma}) = \mathsf{A}\boldsymbol{\Sigma}^{-1}\mathsf{A}^T$ and we defined the functions

$$B_r(h) \equiv \frac{2^{rh}}{\lambda^{rh}(N_s-1)^{rh}}\frac{\Gamma_r\left(\frac{\nu}{2}+h\right)}{\Gamma_r\left(\frac{\nu}{2}\right)}, \tag{35}$$

with $\nu = N_s-1-p+N_p$. Putting $h=1/2$ in Eq. (34) yields the SLI of the mean of $\sqrt{|\mathsf{F}^{-1}_{r\times r}(\mathsf{S})|}$ with respect to a known covariance matrix

$$\mathrm{SLI}_{\mathrm{hypervol.}} = B_r(\frac{1}{2}). \tag{36}$$

In other words, Eq. (36) is the systematically lost information of the r-dimensional hypervolume within ellipsoids of constant credibility. Since the SLI is positive, the parameter volume that is compatible with the data, is inflated, compared to the case of a known covariance matrix. The estimator for the parameter volume as allowed by the true covariance matrix $\boldsymbol{\Sigma}$ is therefore

$$V_\Sigma(r,\mathsf{S}) = \frac{V_e(r,S)}{B_r\left(\frac{1}{2}\right)}, \tag{37}$$

where $V_e(r,S)$ is defined by Eq. (32) using $\mathsf{F}^{-1}(\mathsf{S})$ from Eq. (25). The variance of this debiased parameter volume is then

$$\mathcal{V}(V_\Sigma) = |\mathsf{F}^{-1}_{r\times r}(\boldsymbol{\Sigma})|\frac{[\pi\cdot\chi^2_r(\alpha)]^r}{[\Gamma(\frac{r}{2}+1)]^2}\left[\frac{B_r(1)}{B_r^2(\frac{1}{2})}-1\right]. \tag{38}$$

Specializing to two dimensions, $r=2$, yields the area of credible regions for parameter pairs as displayed in a triangle plot. Eq. (32) is then the area of an ellipse $A=\pi|\boldsymbol{a}_1||\boldsymbol{a}_2|$ where the the semi-axes are given by

$$\boldsymbol{a}_i = \sqrt{\lambda_i\chi^2_2(\alpha)}\,\boldsymbol{e}_i, \tag{39}$$

We can restore the SLI as in Eq. (37) for an ellipse by rescaling its semi-axes:

$$\boldsymbol{a}_i \to \frac{\boldsymbol{a}_i}{\sqrt{B_2(\frac{1}{2})}}. \tag{40}$$

We caution that this rescaling cannot be applied to parameter credible regions of a real dataset. It merely allows a forecast of how much information could be gained, if the covariance matrix were known.

Since $A_T = \pi\chi^2_2(\alpha)\sqrt{|\mathsf{F}^{-1}_{2\times2}(\boldsymbol{\Sigma})|}$ is the area encased by

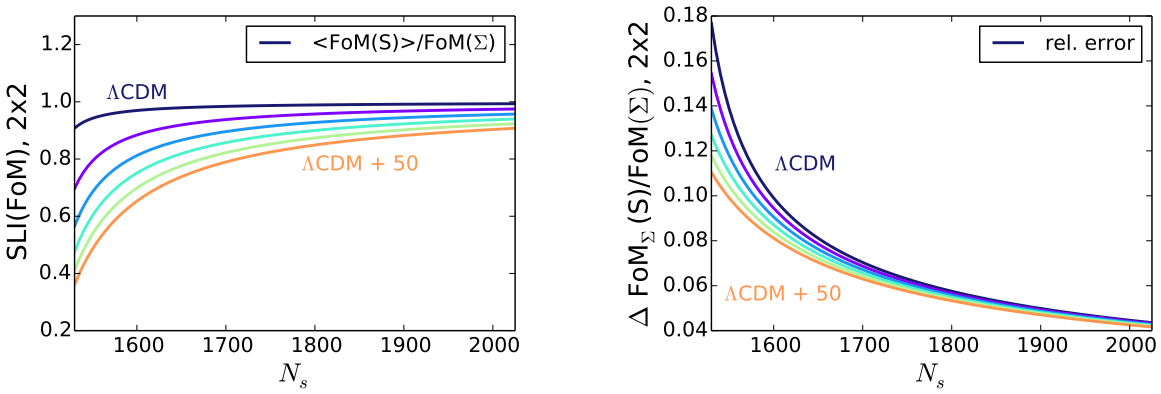

Figure 4. Same Euclid-like setup as in Fig. 3, but now analysing the Figure of Merit (FoM) of 2 parameters, if a total of N_p parameters as indicated by the linestyle are estimated. Left: The SLI decreases if more simulations are run, meaning the figure of merit increases towards its optimal value (which is unity in the plot). Hypothetically restoring the SLI is achieved with Eq. (44) and the right panel depicts the relative error Eq. (46) which describes how accurately the true FoM can be predicted when starting from a an estimator with N_s simulations.

the true credible region, the statistical error of its estimator A_Σ, relative to A_T is

$$\frac{\sqrt{\mathcal{V}(A_\Sigma)}}{\pi\chi^2_2(\alpha)\sqrt{|\mathsf{F}^{-1}_{2\times2}(\mathsf{\Sigma})|}} = \sqrt{\frac{B_2(1)}{B_2^2(\frac{1}{2})} - 1}. \tag{41}$$

The SLI and the relative error for ellipses is given by Eq. (36) for $r = 2$. This SLI is lost in each pair of jointly estimated parameters and is plotted in Fig. 3 for a Euclid-like survey. The left panel demonstrates that the systematically lost information for the credible regions compared to the ideal case of a known covariance matrix. We observe at first a steep decrease of the lost information when more simulations are run, meaning that running only a few more simulations than the minimally possible number, very quickly improves the constraining power. Afterwards, the curves flatten and an asymptotic domain is reached, where the remaining lost information is restored ever more slowly. In contrast, the number of nuisance parameters affects critically how many simulations are needed in order to keep the SLI per parameter-pair below a chosen limit. As the total number of free parameters increases, the SLI increases as well, and a reduction of nuisance parameters can be a more efficient way to reduce the impact of noise in estimated covariance matrices. Fig. 3 shows what could be gained if only more simulations were run, and this analysis will be useful in determining a sweetspot between nuiance parameters and number of simulations when planning future surveys. The right panel depicts the result of restoring the SLI in a forecasting experiment via the estimator A_Σ. This estimator will scatter as depicted. However, its relative error decreases with the number of nuisance parameters as their uncertainty already inflates the size of the credibility regions. Finally, for $r = 1, h = 1$, it follows from Eq. (35) that the SLI in the variance of each single parameter i marginalized over all other parameters is

$$\frac{\langle\sigma_i^2(\mathsf{S})\rangle}{\sigma_i^2(\mathsf{\Sigma})} = \frac{(N_s - 1 - p + N_p)(N_s + 2)}{N_s(N_s - p)}, \tag{42}$$

illustrating that for a given survey with fixed number of simulations N_s and data points p, the lost information scales only linearly with the number of nuisance parameters.

4 SLI IN THE FIGURE OF MERIT

The inverse of the parameter volume that is compatible with the data, is a useful figure of merit. In the Gaussian approximation, this volume is proportional to the square root of the determinant of the inverse Fisher matrix, where the inverse is taken, because ignoring rows and columns of it allows to quickly marginalize parameters. We therefore work with the figure of merit

$$\mathrm{FoM} = \frac{1}{\sqrt{|\mathsf{F}^{-1}_{r\times r}(\mathsf{S})|}}. \tag{43}$$

Eq. (43) is the FoM conditional on S, and has lost information due to the finite number of simulations in S. The mean lost information with respect to the case of a known $\mathsf{\Sigma}$ follows from putting $h = -1/2$ in Eq. (34). Dividing by the determinant of the true Fisher matrix yields the following estimator for the figure of merit that restores the SLI

$$\mathrm{FoM}_\Sigma(\mathsf{S}) = \frac{1}{\sqrt{|\mathsf{F}^{-1}_{r\times r}(\mathsf{S})|}}\frac{1}{B_r(-\frac{1}{2})}. \tag{44}$$

If $r = N_p$, Eq. (43) describes the figure of merit for the entire parameter set, if $r = 2$, Eq. (43) allows to define a figure of merit e.g. for the w_0, w_a subspace of a generic dark energy model. The variance of the estimated figure of merit with the restored SLI is then

$$\mathcal{V}[\mathrm{FoM}_\Sigma(\mathsf{S})] = [\mathrm{FoM}(\mathsf{\Sigma})]^2\left[\frac{B_r(-1)}{B_r^2(-\frac{1}{2})} - 1\right], \tag{45}$$

resulting in the relative error with respect to the true figure of merit

$$\frac{\sqrt{\mathcal{V}[\mathrm{FoM}_\Sigma(\mathsf{S})]}}{\mathrm{FoM}(\mathsf{\Sigma})} = \sqrt{\frac{B_r(-1)}{B_r^2(-\frac{1}{2})} - 1}. \tag{46}$$

The SLI and the remaining relative error after restoring the SLI are depicted in Fig. 4.

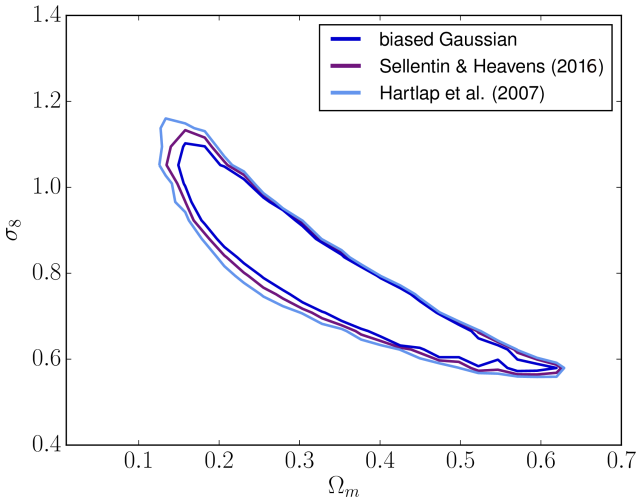

Figure 5. 68% credible regions for the matter density parameter Ω_m and the amplitude of the linear perturbations σ_8, derived from the Hartlap et al. likelihood (outer), from the Sellentin-Heavens likelihood (middle), and from the ideal Gaussian likelihood (inner), which deliberately ignores uncertainties in the estimated covariance matrix, in order to illustrate that the dominant fraction of the total parameter uncertainty arises from other effects than covariance uncertainty.

5 APPLICATION TO CURRENT SURVEYS

5.1 KiDS and DES

The KiDS-450 survey (Hildebrandt et al. 2016) employs $p = 130$ data points, and a covariance matrix estimated from 930 simulations. Estimating about 10 cosmological parameters, the parameter credibility contours are inflated at the insignificant level of 1% or less. The statistical deviations of the contours from their true position is then about 3%, allowing to infer that KiDS does not suffer from noise in the covariance matrix.

The DES science verification (SV) weak lensing analysis uses $p = 36$ data points and a covariance matrix estimated from 126 simulations. For estimating about 10 cosmological parameters, we therefore detect a 10% SLI in the DES 2D marginal contours, together with an additional 10% statistical scatter. DES would need to increase the number of simulations to ≈ 450 in order to reduce the lost information to 5%, keeping the dimension of the data vector fixed to $p = 36$. For the analysis of the full DES data set, the number of simulations needed can be estimated with Eq. (37).

Since all of the previous results hold in the Fisher matrix approximation, we now check on their validity by running a full likelihood, comparing with our expectations. We choose the weak lensing measurement of the DES SV data, varying all parameters of ΛCDM together with two intrinsic alignment parameters. Physically, our setup is exactly that presented in Abbott et al. (2015), and we modify the statistical analysis only. Fig. 5 depicts our results. The total uncertainty on the combination $\sigma_8 - \Omega_m$ includes the degeneracy of the parameters in the ΛCDM model, the finite number of data points and their noise, and finally the uncertainty of the covariance matrix. Deliberately ignoring all statistical uncertainty in the covariance matrix, allows us to estimate how much of the error bars on σ_8 and Ω_m is due to the parameter degeneracies and the finiteness of the noisy data set. This produces the innermost blue contour seen in Fig. 5, which gives a good estimate of the optimally achievable parameter bounds. Including the uncertainty of S by using the t-distribution produces the middle contour in Fig. 5. For comparison, we also plot the contour derived with the procedure proposed by Hartlap et al. (2007), which produces the outermost contour. The closeness of all contours in Fig. 5, reveals that the primary source of parameter uncertainty is not the uncertainty of the covariance matrix, but rather the degeneracy of the parameters and the finiteness of the data set. This finding is in line with our findings in Fig. 3 and Fig. 4 where we found for Euclid that the number of nuisance parameters is the more pressing issue than the number of simulations for the covariance matrix. We furthermore see that our Fisher matrix predictions capture the trend of the credibility contours well: Updating from Hartlap et al.'s method to the t-distribution moves the contours in. Note that although the corrections seem small, the tightening of the credible regions is roughly equivalent to observing half a million (corresponding to 20%) more galaxies. We made sure none of these results are due to numerical issues with running Monte-Carlo Markov Chains: multiple chains were run, which fulfil the Gelman-Rubin convergence test. The exchange of the statistical analysis was achieved by reweighting the chains, instead of resampling, such that all differences in the contours come from the different statistical likelihoods, rather than from different sampling.

5.1.1 *Flat vs non-flat Universe*

Finally, we note that one might expect the broader tails of the Sellentin & Heavens (2016) likelihood to have a significant effect on Bayesian evidence calculations. We have tested this with the DES SV data, using nested sampling (Feroz et al. 2009) within cosmosis (Zuntz et al. 2014) to compute the Bayes factor for flat and non-flat Universes, with a prior range ± 0.2 for the curvature parameter Ω_k. Both the Sellentin & Heavens (2016) and Hartlap et al. (2007) likelihood formally favour a non-flat Universe ($\ln(B) = 0.17 \pm 0.09$ and 0.30 ± 0.07 respectively), but neither significantly, and the difference in Bayes factors is insignificant in this case, but the sign of the change is as expected.

In the context of model comparison, we note that the r-dimensional Figure of Merit (and its associated parameter volume V_r) calculated earlier plays an interesting role, closely related to Lindley's paradox. Consider two nested models, in which the simpler model A corresponds to a special point in B, setting r parameters in B to specific values. If the data have low probability in A, model A may still be preferred in a Bayesian analysis, provided that the prior probability that the parameters are close to the special point is even smaller in B. The Bayes factor essentially compares these two potentially small probabilities. It is roughly the ratio of $\exp(-\chi^2/2)$ in model A to the probability in B that the special point lies within the volume V_r when the centre of V_r is chosen randomly from the prior volume. For a uniform prior, this latter probability is essentially V_r divided by the prior volume.

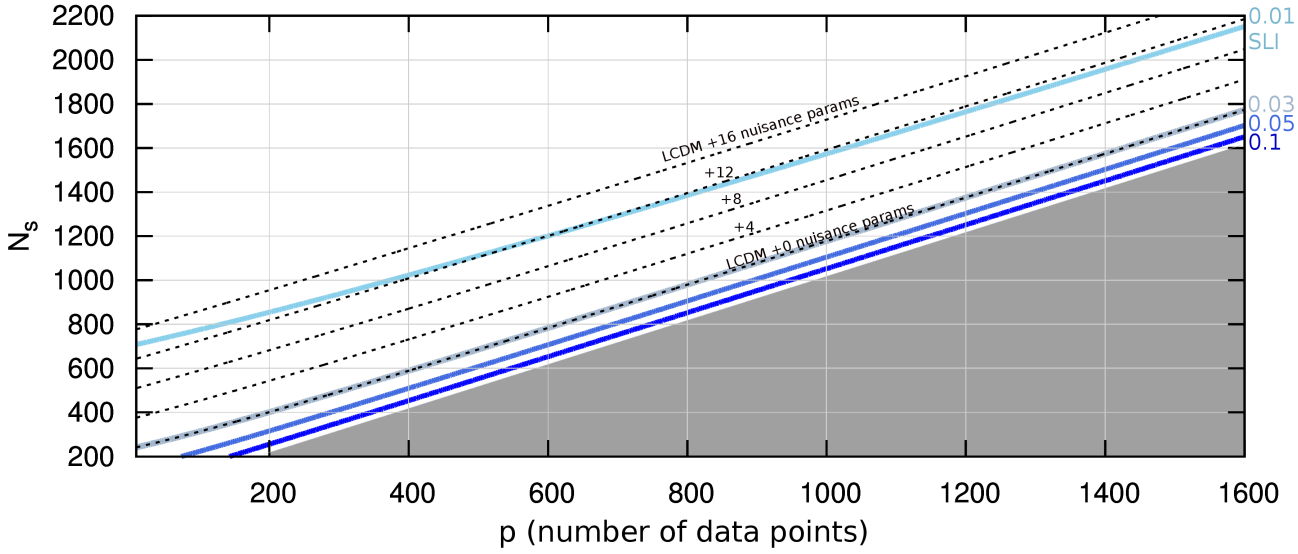

Figure 6. Orientation for future surveys: the number of simulations N_s needed to reach a certain precision, given the number of data points p and the number of parameters. The grey triangle is excluded, because at least as many simulations as data points have to be run. Its edge is the minimal number of simulations that need to be run. The thick blue lines mark a 6-parameter ΛCDM model without any nuisance parameters, labeled in blue hues by the different percentages of systematically lost information. The dashed black lines all refer to a 3% loss of information, with an additional number of nuisance parameters as indicated in black. The quantity here plotted is the lost information in the variance of single parameters marginalized over all other jointly estimated parameters, see Eq. (42). The needed N_s for reaching a comparable precision in joint parameter confidence contours, or in model comparison, is then slightly higher.

6 CONCLUSIONS

If the data covariance matrix in parameter inference is estimated from N_s simulations, its uncertainty leads to a loss of information on cosmological parameters. In this paper, we have quantified the systematically lost information (SLI) with a Fisher matrix approach starting from the t-distribution Eq. (2) found in Sellentin & Heavens (2016), which is the Bayesian solution to propagating the covariance matrix uncertainty into the parameters.

Within our approximations, the SLI corresponds to a 'bias' of the t-distributions credible regions with respect to the narrower credible regions derived from a true covariance matrix. This lost information can be reduced in a real data analysis by increasing the number of simulations. Due to Jensen's inequality, this SLI in the Figure of Merit and joint parameter credibility regions cannot be assessed by removing a bias in the Fisher matrix or its inverse, which is the level at which previous works have stopped (Dodelson & Schneider 2013; Taylor & Joachimi 2014). We therefore computed the SLI of these quantities directly, see Sects. 4 -5. This allows us to estimate how far away a given survey is from its optimally achievable error bounds.

We summarize our scientific conclusions as follows. We find that for current weak lensing surveys, as KiDS-450 and DES SV, the parameter degeneracies and the finite data set are the dominating drivers of the uncertainties of cosmological parameters; uncertainty due to the covariance matrix is subdominant. KiDS systematically lost about 1% of information, the analysis of the DES SV about 10%. The exact numbers depend on the quantity for which the information loss is estimated, e.g. the variance of a single parameter, the credibility region of all jointly estimated parameter pairs, the entire parameter volume, or the Figure of Merit.

Concerning a Euclid-like weak lensing survey with $p = 1500$ data points, Eq. (37) predicts that in a pure ΛCDM measurement Euclid would systematically lose only 1% of information for 1900 simulations, as compared to knowing the covariance matrix. This is only about twice as many simulations as have been run for KiDS. An estimation of the optimally achievable error bound can then be achieved with a 4% uncertainty. In the Bayesian framework of evaluating the t-distribution, this 4% uncertainty does not need to be added to the parameter uncertainties anymore – the t-distribution already accounts for this. The uncertainty merely refers to the reliability of predicting the optimally achievable error bound. For 10 additional nuisance parameters 2900 simulations are sufficient to only lose 1% of systematic uncertainty, leading to a remaining scatter of 2%.

In general, the Wishart nature of estimated covariance matrices already enforces that at least as many simulations as the dimension of the data set are run. We found for all estimators here considered, that a minor increase of the simulations over this minimally needed number already limits the lost information to below 10% or 5%. Afterwards, an asymptotic domain is reached, where a further increase of N_s leads to an ever slower restoration of the remaining lost information. Fortunately, reducing the number of nuisance parameters reduces the lost information due to covariance matrix estimation much more quickly. In the case of fully marginalized variances of individual parameters, the SLI reduces linearly with a reduction of nuisance parameters. For all weak lensing surveys, this highlights the desirability of obtaining a better physical understanding of intrinsic alignments and shape measurement challenges, so that only a limited number of nuisance parameters are required.

For future surveys, we compress the findings in this pa-

per into Fig. 6, which displays the needed number of simulations to reach a certain precision, given a targeted number of data points, and a known number of cosmological and nuisance parameters, as indicated by the line style.

In comparison to previous forecasts (Dodelson & Schneider 2013; Taylor et al. 2013; Taylor & Joachimi 2014; Percival et al. 2014), we find here significantly lower figures for the needed number of simulations in order to aquire a certain precision in parameter inference. This is mainly due to our use of the t-distribution, Eq. (2) which is more accurate than the scaled Gaussian likelihood as proposed by (Hartlap et al. 2007) and employed in Dodelson & Schneider (2013); Taylor et al. (2013); Taylor & Joachimi (2014); Percival et al. (2014).

As a side effect, this study has produced all quantities needed for forecasts with estimated covariance matrices. Eq. (17) gives the Fisher matrix, as e.g. needed for Euclid.

The optimally achievable error bounds can then be obtained using the rescaling of ellipse-axes as in Eq. (40).

7 ACKNOWLEDGEMENTS

We thank Joe Zuntz, Sarah Bridle and Andrew Jaffe for discussions, support and comments. cosmosis can be found at https://bitbucket.org/joezuntz/cosmosis/wiki/Home. This work was supported by a postdoc fellowship of the German Academic Exchange Service (DAAD), awarded to ES.

References

Abbott T., Abdalla F. B., Allam S., Amara A., Annis J., Armstrong R., Bacon D., Banerji M., Bauer A. H., Baxter E., Becker M. R., Bernstein R. A., 2015, ArXiv e-prints, 1507.05552

Anderson T. W., 2003, An Introduction to Multivariate Statistical Analysis. Wiley Series in Probability and Statistics, 3rd ed.

Blot L., Corasaniti P. S., Alimi J.-M., Reverdy V., Rasera Y., 2015, MNRAS, 446, 1756

Dodelson S., Schneider M. D., 2013, Phys. Rev. D, 88, 063537

Feroz F., Hobson M. P., Bridges M., 2009, MNRAS, 398, 1601

Gupta A. K., Nagar D. K., 2000, Matrix variate distributions. Chapman & Hall/CRC

Harnois-Déraps J., van Waerbeke L., 2015, MNRAS, 450, 2857

Hartlap J., Simon P., Schneider P., 2007, A,A, 464, 399

Heymans C., et al. 2013, MNRAS, 432, 2433

Hildebrandt H., Viola M., Heymans C., Joudaki S., Kuijken K., Blake C., Erben T., Joachimi B., Klaes D., Miller L., Morrison C. B., Nakajima R., Verdoes Kleijn G., 2016, ArXiv e-prints, 1606.05338

Jain B., Spergel D., Bean R., Connolly A., Dell'antonio I., Frieman J., Gawiser E., Gehrels N., Gladney L., 2015, ArXiv e-prints, 1501.07897

Laureijs R., Amiaux J., Arduini S., Augueres J. ., Brinchmann J., Cole R., Cropper M., Dabin C., Duvet L., Ealet A., et al. 2011, ArXiv e-prints, 1110.3193

Percival et al. W. J. e., 2014, MNRAS, 439, 2531

Sato M., Takada M., Hamana T., Matsubara T., 2011, APJ, 734, 76

Sellentin E., Heavens A. F., 2016, MNRAS, 456, L132

Taylor A., Joachimi B., 2014, MNRAS, 442, 2728

Taylor A., Joachimi B., Kitching T., 2013, MNRAS, 432, 1928

Tegmark M., Taylor A., Heavens A., 1997, Astrophys.J., 480, 22

Zuntz J., Paterno M., Jennings E., Rudd D., Manzotti A., Dodelson S., Bridle S., Sehrish S., Kowalkowski J., 2014, ArXiv eprints, 1409.012

This paper has been typeset from a TEX/ LATEX file prepared by the author.